CS 294: Big Data System Research: Trends and Challenges

Fall 2015 (MW 9:30-11:00, **310 Soda Hall**)

Ion Stoica and Ali Ghodsi

(http://www.cs.berkeley.edu/~istoica/classes/cs294/15/)

Big Data

First papers:

- 2003: The Google file system paper
- 2004: The MapReduce paper

Today every major system & networking conference has Big Data sessions

Big Data Impact

Already helped create new business

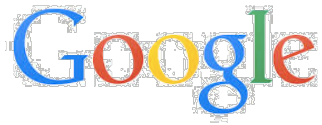

Already helped disrupt existing businesses

- » Retail 

- » Rental 

- » Taxi 

- » home appliances 

- » …

Big Data Stack

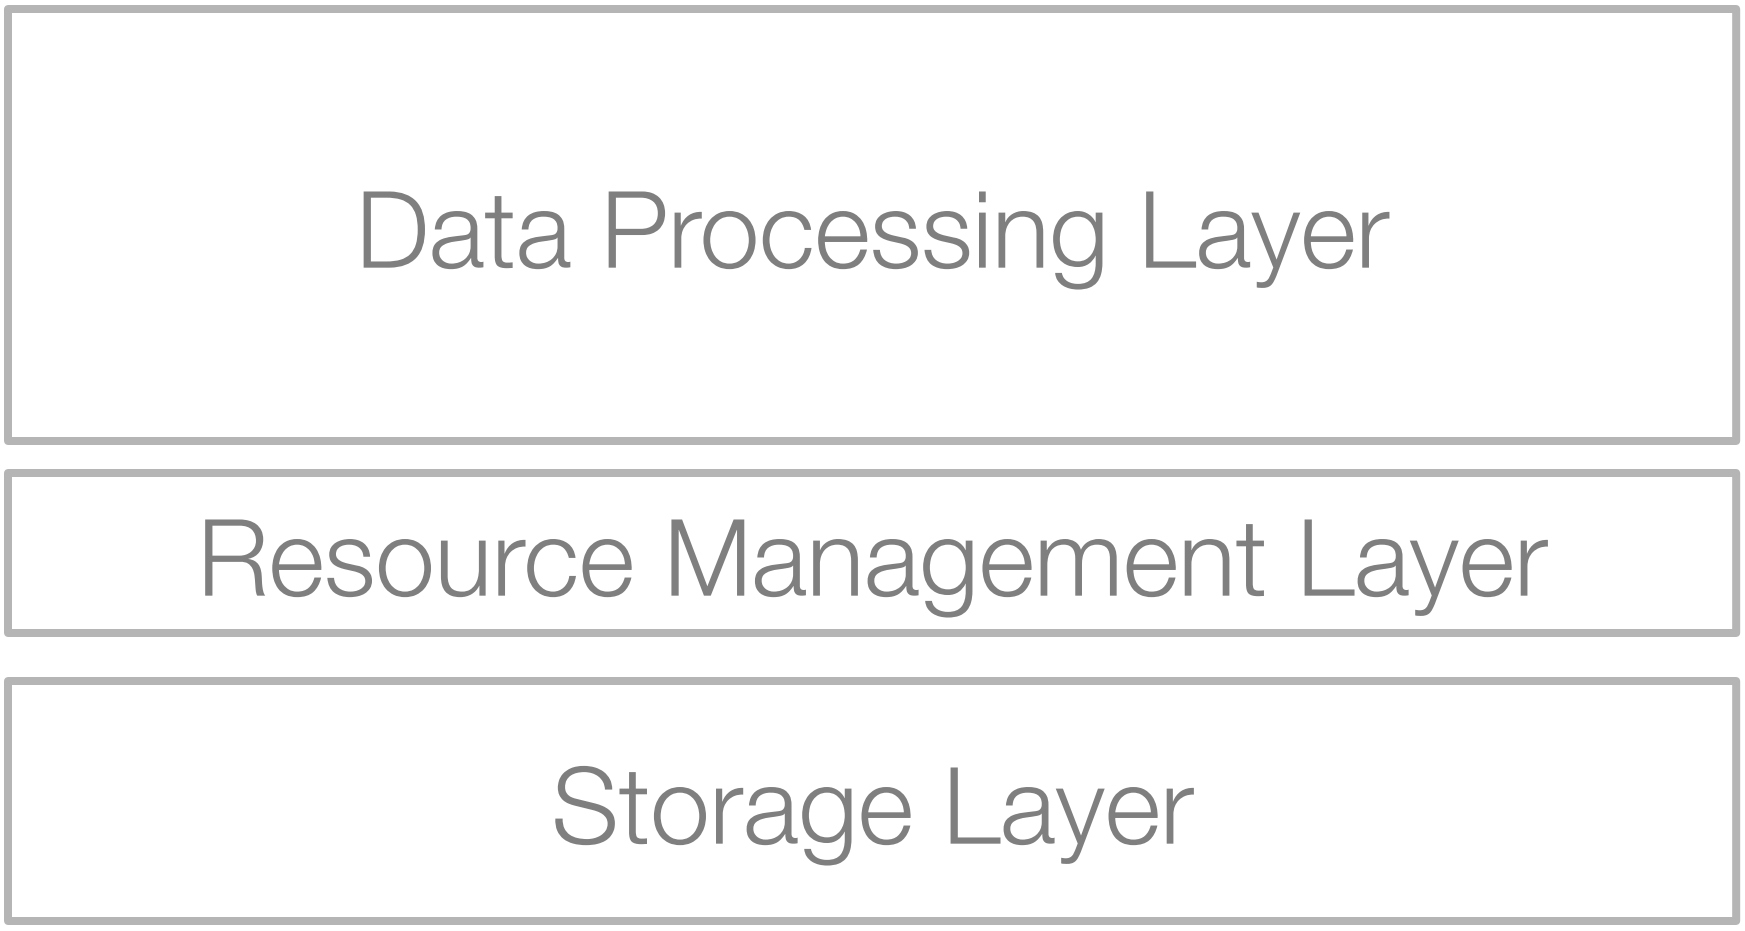

Hadoop Stack

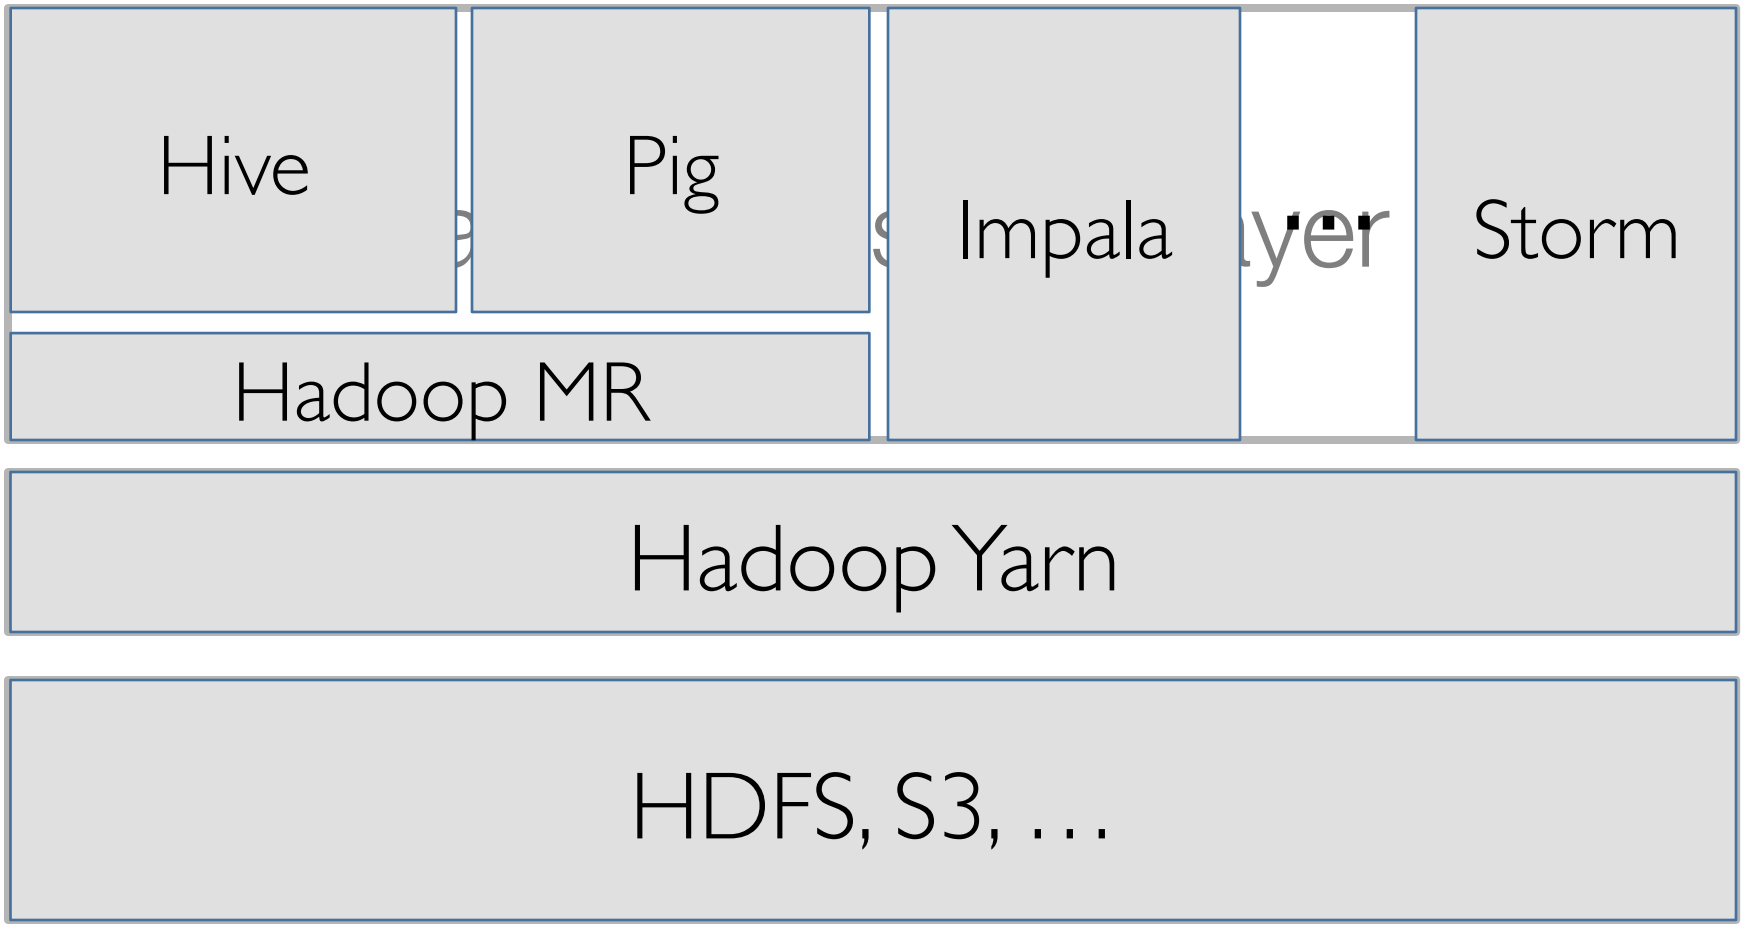

The Berkeley AMPLab

January 2011 – 2017
- » 8 faculty
- » > 40 students
- » 3 software engineer team

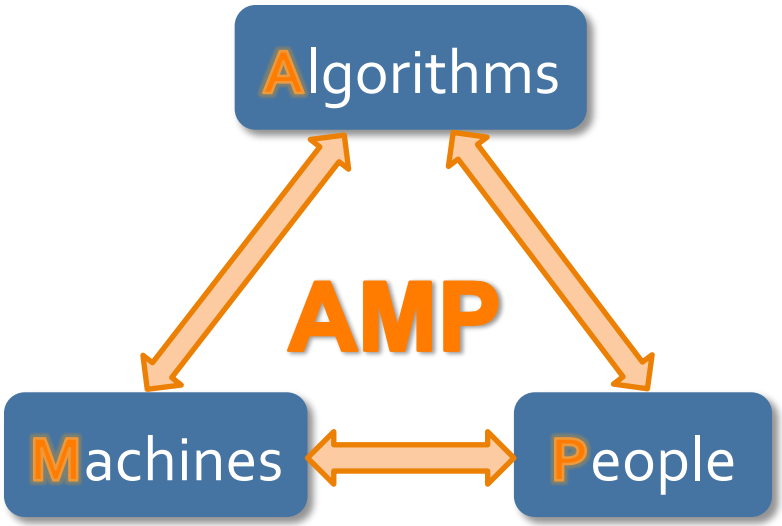

Organized for collaboration

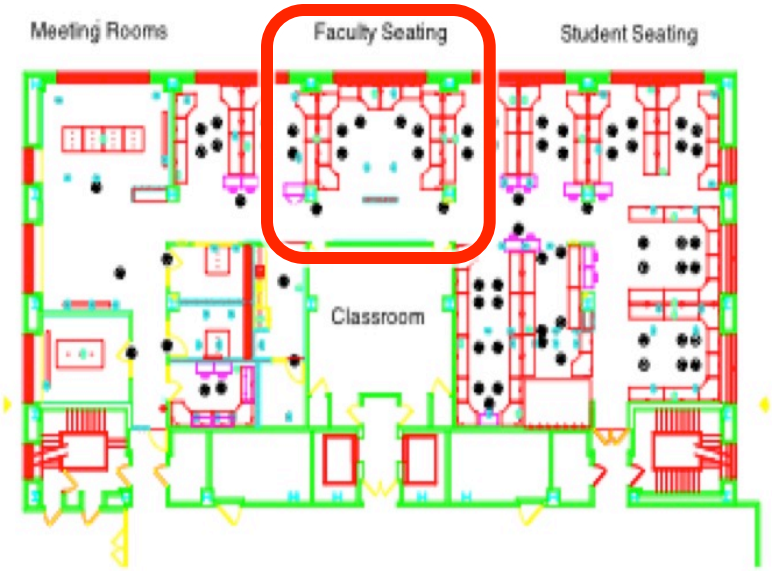

3 day retreats
(twice a year)

AMPCamp3
(August, 2013)

220 campers
(100+ companies)

The Berkeley AMPLab

Governmental and industrial funding:

Goal: Next generation of open source data analytics stack for industry & academia:
Berkeley Data Analytics Stack (BDAS)

BDAS Stack

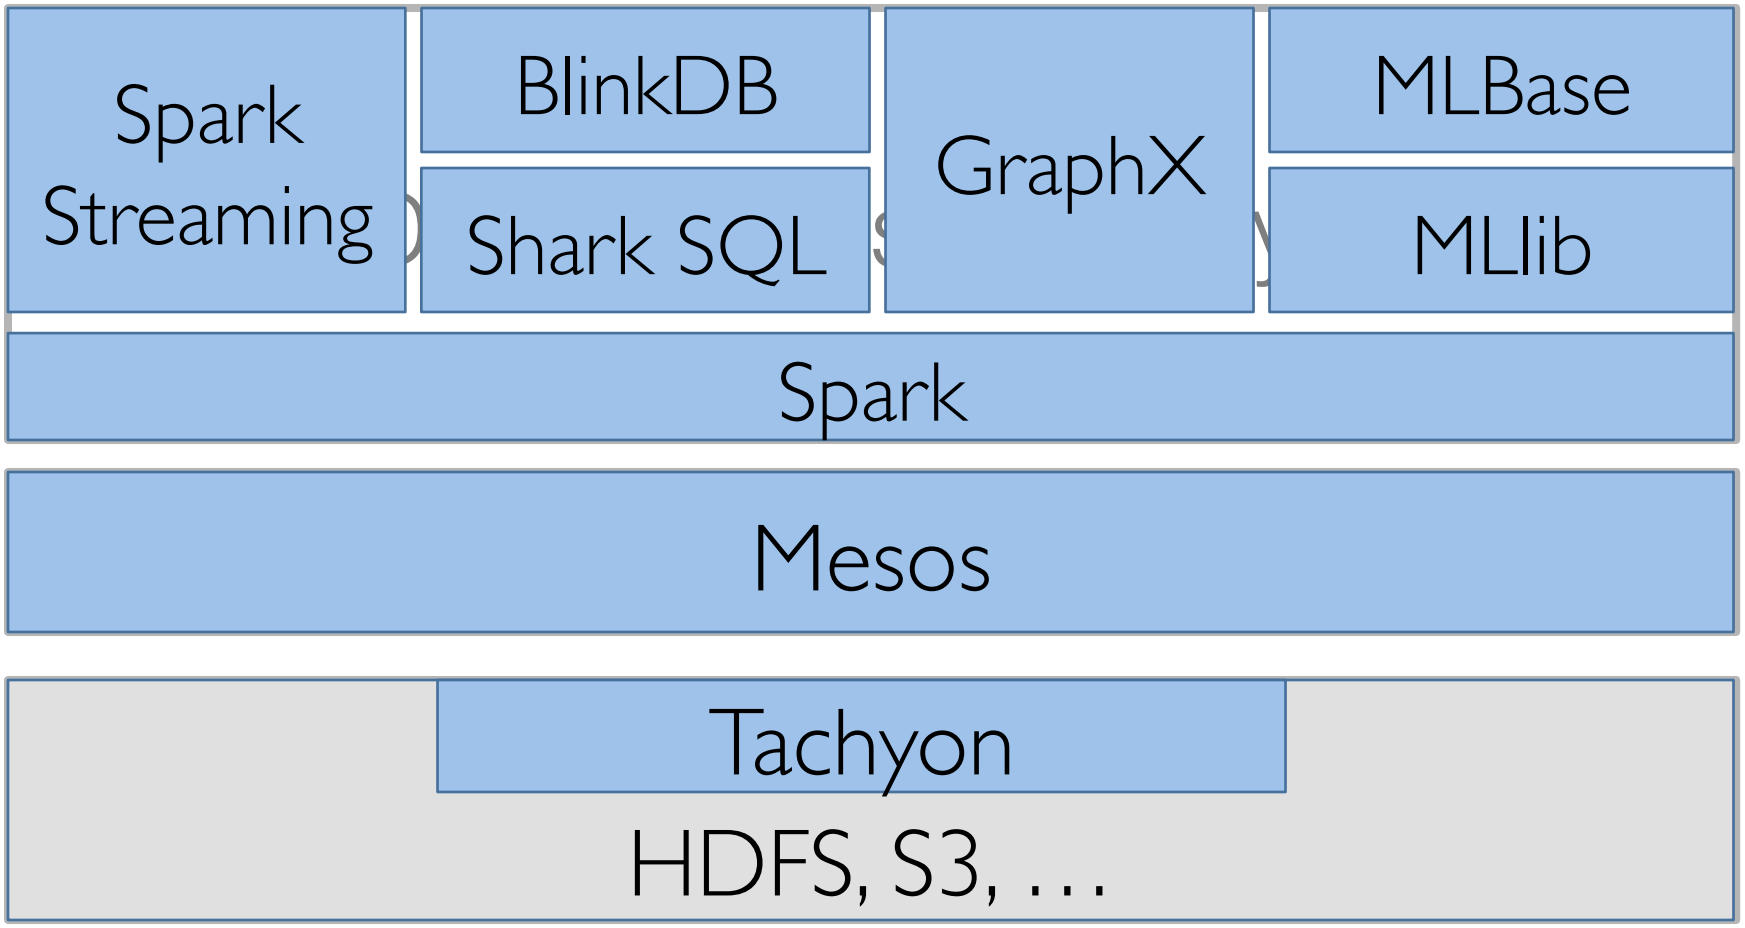
Spark Streaming
BlinkDB
Shark SQL
GraphX
MLBase
MLlib
Spark
Mesos
Tachyon
HDFS, S3, …

BDAS & Hadoop fitting together

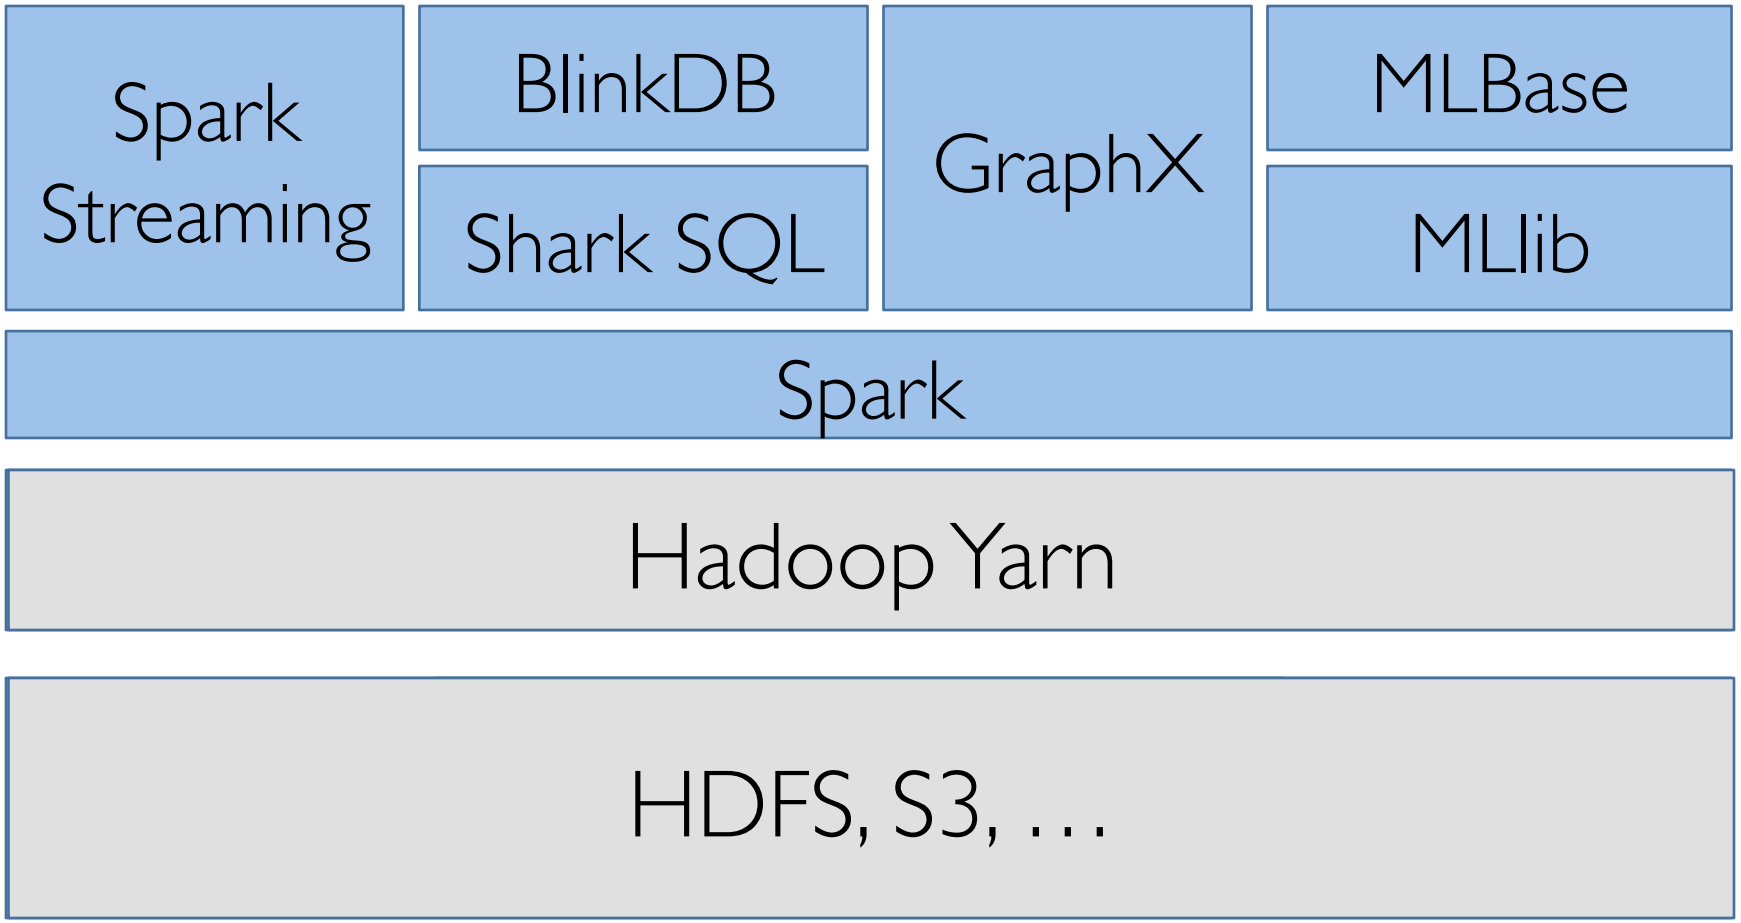

How do BDAS & Hadoop fit together?

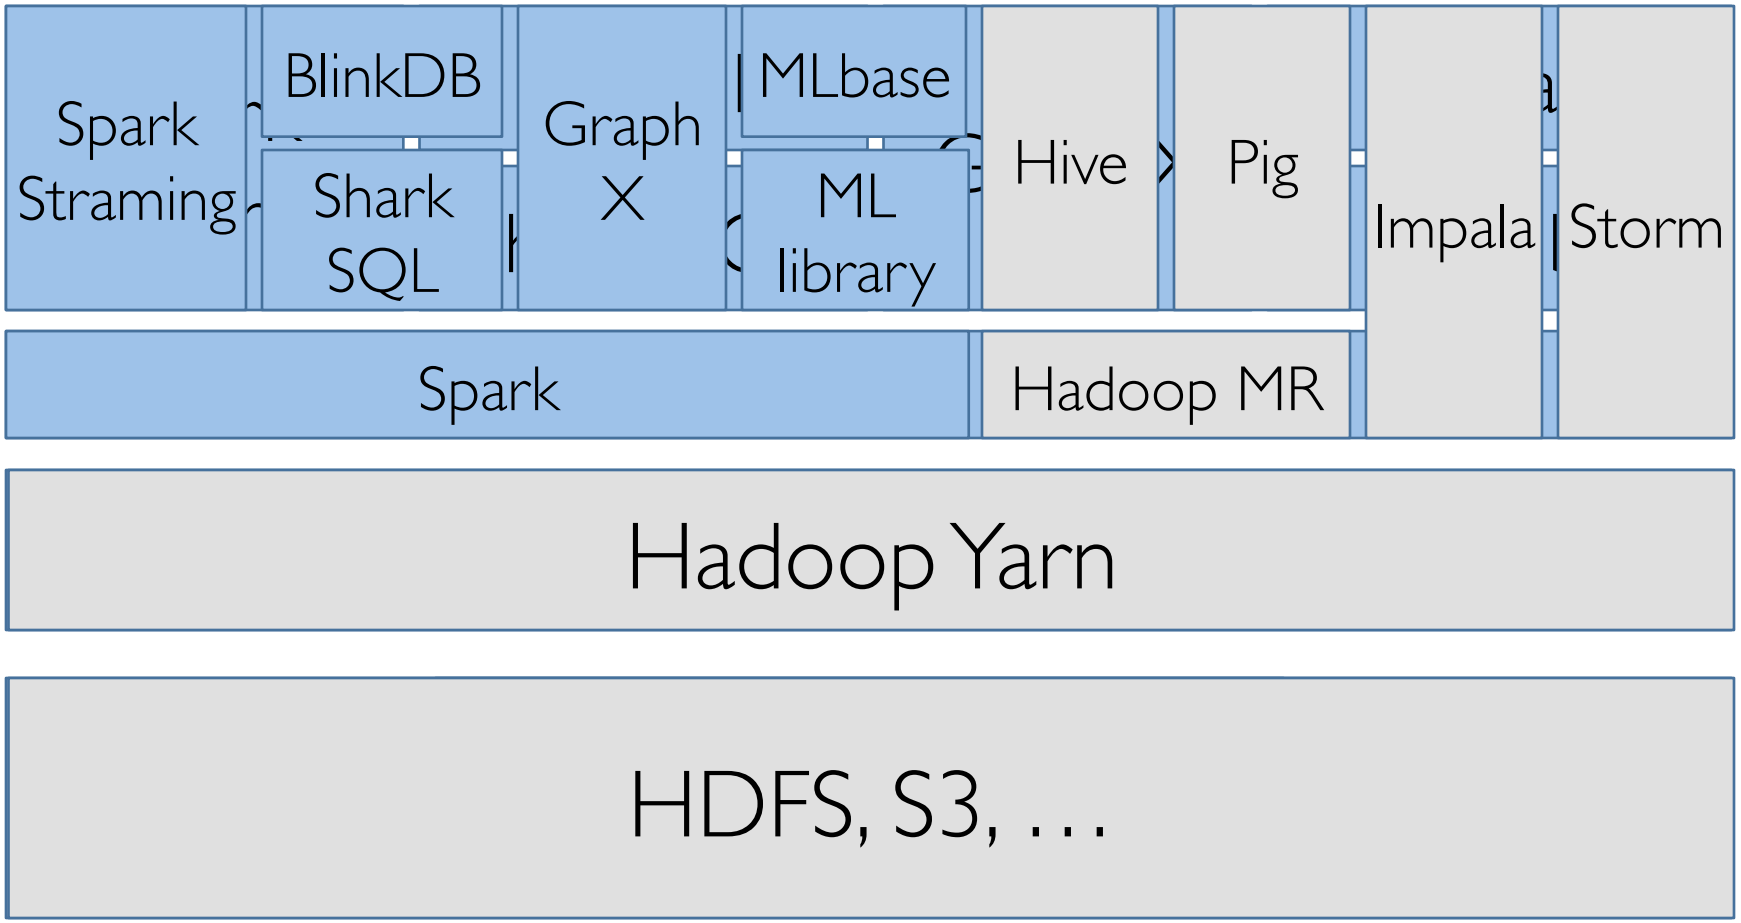

How do BDAS & Hadoop fit together?

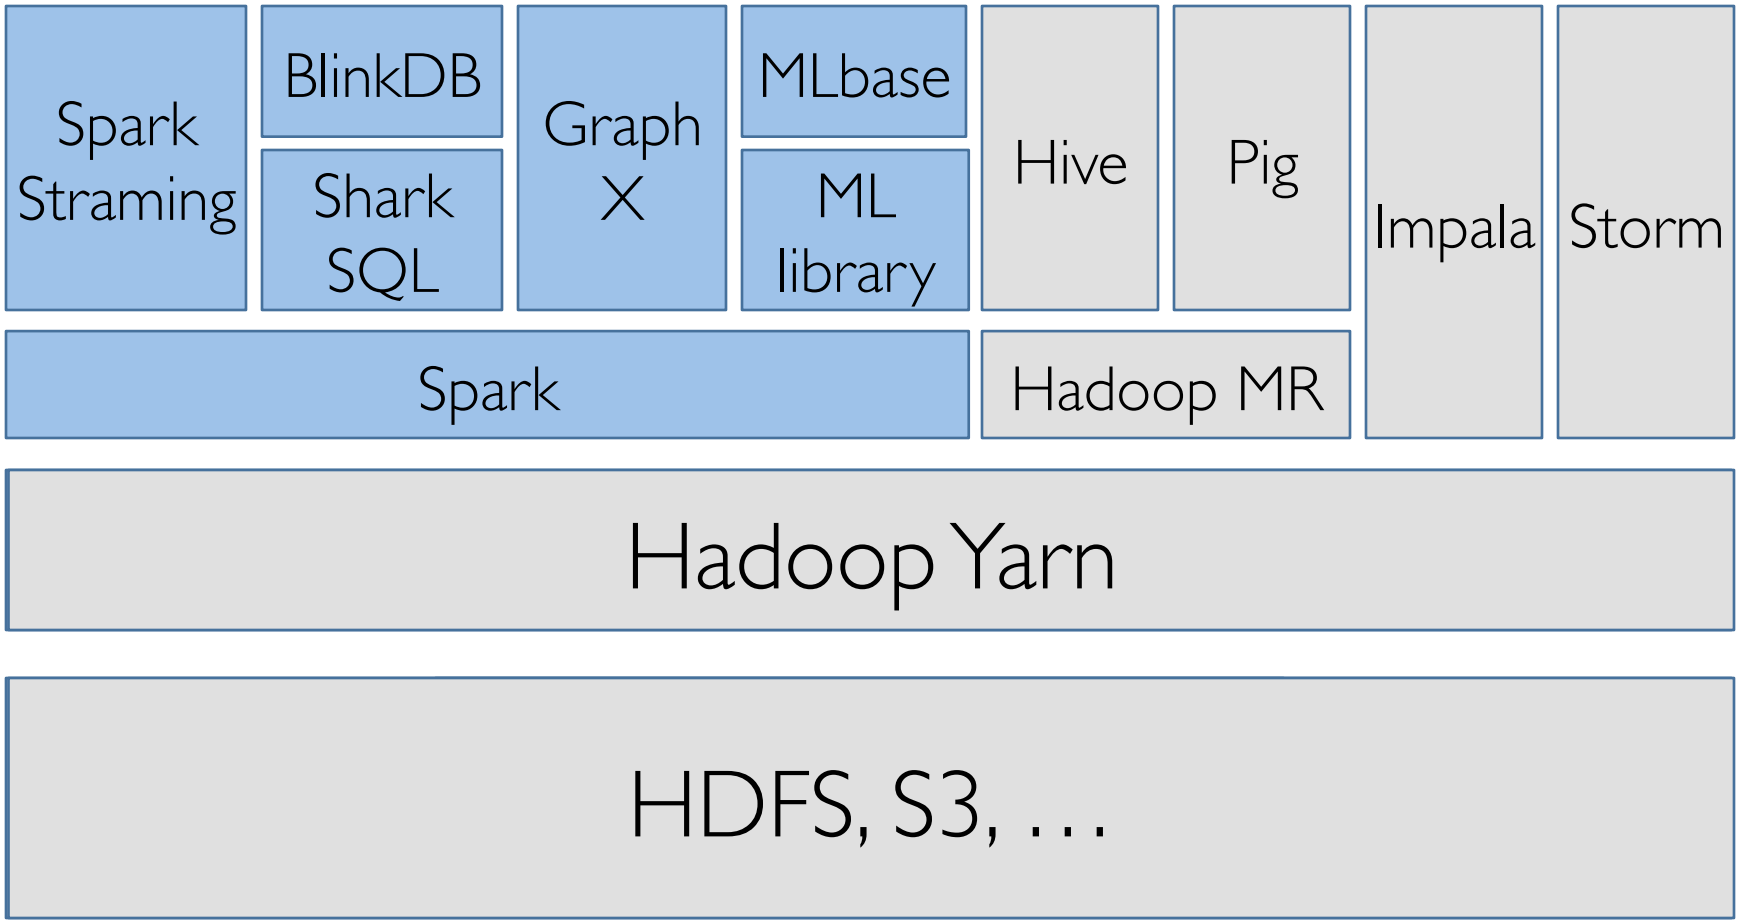

This Class

Learn about state-of-art research in Big Data

Work on an exciting project

Hopefully start next generation of impactful projects

Grading

Project: 60%

Class presentations: 40%

- Around 2 papers per student
- See Randy's guidelines for leading discussion on papers
  - http://bnrg.eecs.berkeley.edu/~randy/Courses/CS294.F07/LeadingPapers.pdf

Administrative Information

Class website:
http://www.cs.berkeley.edu/~istoica/classes/cs294/15/

Office Hours (Soda 465D):
- » TBA

Create an (anonymized) blog account for paper reviews if you don't have one yet (e.g., www.blogger.com)
- » Sent me an e-mail by Monday, August 31, with your blog url
- » Preferred e-mail for the class e-mail list

Papers

Is the problem real?

What is the solution's main idea (nugget)?

Why is solution different from previous work?

- » Are system assumptions different?
- » Is workload different?
- » Is problem new?

Does the paper (or do you) identify any fundamental/hard trade-offs?

Papers (cont'd)

Do you think the work will be influential in 10 years?

- Why or why not?

Predicting the future hard, but worth a try

- Look at past examples for inspiration

Streaming Over TCP

Countless papers:

- » Why cannot be done…
- » New protocols to do it…

Today

- » Virtually all streaming over TCP
- » Trend to stream over HTTP!

Why did it Succeed?

Multicast

Countless papers:

- » Why world will come to a standstill without multicast…
- » New protocols to do it…

Today

- » Multicast is used only in enterprise settings at best
- » Overlay multicast widely used in the Internet
  - CDN based, e.g., WorldCup, March Madness, Iinagurations, ...
  - P2P, mostly popular outside US (e.g., China)

Why Did it Fail?

Shared Memory

Countless papers:

- How shared memory simplifies programming parallel computers
- Many, many systems proposed and build

Today:

- Message passing (MPI) took over as the de facto standard for writing parallel applications

Why Did it Fail?

Network Computer

Big in 90s

- Promoted by an alliance of Sun, Oracle, Acorn

Promise: many of advantages of cloud computing

- Easy to manage
- Application sharing
- …

Failed miserably

Why Did it Fail?

Coming Back: ChromeOS

Will it succeed this time?

What are Hard/Fundamental Tradeoffs?

Brewer's CAP conjecture: "Consistency, Availability, Partition-tolerance", you can have only two in a distributed system

In a *in-order, reliable* communication protocol cannot minimize overhead and latency simultaneously

Hard to simultaneously maximize evolvability and performance